S-PowerGraph: Streaming Graph Partitioning for Natural Graphs by Vertex-Cut

Cong Xie
Department of Computer Science and Engineering
Shanghai Jiao Tong University
800 Dong Chuan Road
Shanghai, China 200240
xcgoner1108@gmail.com

Wu-Jun Li
National Key Laboratory for Novel Software Technology
Department of Computer Science and Technology
Nanjing University
liwujun@nju.edu.cn

Zhihua Zhang
Department of Computer Science and Engineering
Shanghai Jiao Tong University
800 Dong Chuan Road
Shanghai, China 200240
zhang-zh@cs.sjtu.edu.cn

arXiv:1511.02586v1 [cs.SI] 9 Nov 2015

ABSTRACT

One standard solution for analyzing large natural graphs is to adopt distributed computation on clusters. In distributed computation, graph partitioning (GP) methods assign the vertices or edges of a graph to different machines in a balanced way so that some distributed algorithms can be adapted for. Most of traditional GP methods are *offline*, which means that the whole graph has been observed before partitioning. However, the offline methods often incur high computation cost. Hence, *streaming graph partitioning* (SGP) methods, which can partition graphs in an online way, have recently attracted great attention in distributed computation. There exist two typical GP strategies: edge-cut and vertex-cut. Most SGP methods adopt edge-cut, but few vertex-cut methods have been proposed for SGP. However, the vertex-cut strategy would be a better choice than the edge-cut strategy because the degree of a natural graph in general follows a highly skewed power-law distribution. Thus, we propose a novel method, called *S-PowerGraph*, for SGP of natural graphs by vertex-cut. Our S-PowerGraph method is simple but effective. Experiments on several large natural graphs and synthetic graphs show that our S-PowerGraph can outperform the state-of-the-art baselines.

Categories and Subject Descriptors

A.3 [**Distributed computing-cloud, map-reduce, MPI, others**]: Big data; 7 [**Data streams**]; 13 [**Graph mining**]

General Terms

Algorithms, Experimentation

Keywords

Graph Partitioning, Streaming, Vertex-Cut

1. INTRODUCTION

In this big data era, data mining and machine learning with massive datasets have captured more and more attentions from research community and industry. In particular, recent years have witnessed the emergence of large natural graphs in many applications like Twitter and Facebook. It is very difficult for a single commodity machine to handle such kinds of large-scale graphs. Hence, more and more works try to adopt parallel (distributed) computation on clusters for tasks with large graphs. To complete distributed computation, we first need to partition the whole graph across the cluster. Hence, *graph partitioning* (GP) [14, 5, 6, 13] has become one of the hot topics in graph-related research.

In distributed computation with clusters, the inter-machine communication cost is typically high. Hence, the goal of GP is to minimize the inter-machine (cross partition) communication, and simultaneously keep the number of vertices (or edges) in every partition approximately balanced [14]. There exist two representative GP strategies: one is edge-cut and the other is vertex-cut. Edge-cut [14, 12, 15] divides the graphs by cutting the edges and assigns the vertices to different partitions. To achieve the goal of GP, edge-cut tries to minimize the cross-partition edges and keep the number of vertices in every partition balanced. On the contrary, vertex-cut [5, 6] divides the graphs by cutting the vertices and assigns the edges to different partitions. To achieve the goal of GP, vertex-cut tries to minimize the number of replications of vertices and keep the number of edges in every partition balanced. Figure 1 illustrates the difference between these two strategies. In Figure 1 (a), the edges (A,C) and (A,E) are cut. The two partitions are node sets {A,B,D} and {C, E}, respectively. In Figure 1 (b), the node A is cut. The two partitions are edge sets {(A,B), (A,D)} and {(A,C), (A,E)}, respectively. It is easy to see that there exists essential difference between these two strategies.

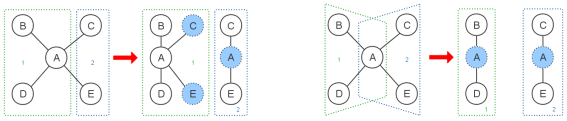

(a) Edge-Cut (b) Vertex-Cut

Figure 1: Two strategies for graph partitioning. Shaded vertices are ghosts and mirrors respectively.

Most of the traditional GP methods are *offline* [7], which means that the whole graph has been observed before partitioning. However, the offline methods often incur high computation cost and hence can not scale to big data applications. Furthermore, the system often need to serially read the graph data from some disks onto a cluster or dynamically crawl the graph from some web sites. Hence, *streaming graph partitioning* (SGP) methods [14, 6, 12, 15], which can partition graphs in an online way, have recently attracted more and more attention from researchers. Representative methods include linear deterministic greedy (LDG) [14] and FENNEL [15]. However, most existing SGP methods adopt the edge-cut strategy for partitioning.

Actually, many existing offline (non-streaming) works, such as PowerGraph [5], have found that the vertex-cut strategy can outperform edge-cut in applications with natural graphs which in general have highly skewed power-law degree distributions. Even in the edge-cut based SGP work [14], the authors admitted that the *edge-balanced partition*, which can be achieved by vertex-cut, may be more preferable than edge-cut for power-law distributed graphs. However, few vertex-cut methods have been proposed to solve the SGP problem.

In this paper, a novel method called *S-PowerGraph*[1] is proposed for SGP of natural graphs by vertex-cut. The main contributions of this paper are briefly outlined as follows:

- This is the first work to systematically study and compare different heuristics for SGP of natural graphs by vertex-cut.
- Although GraphBuilder [6] and PowerGraph [5] are not motivated by SGP and hence cannot be directly used for SGP, we propose methods to adapt them for SGP.
- S-PowerGraph can be viewed as an improved version of the *adapted PowerGraph*, by making use of the property of power-law degree distribution for vertex-cut based SGP.
- Experiments on several large natural graphs and synthetic graphs show that our S-PowerGraph can outperform other baselines, including the adapted GraphBuilder and PowerGraph, to achieve the state-of-the-art performance in real applications.

2. PROBLEM FORMULATION

Graph G is denoted by $G = (V, E)$ where $V = \{v_1, v_2, \ldots, v_n\}$ is a set of vertices and $E \subseteq V \times V$ denotes a set of edges. This implies that the number of vertices $|V| = n$, i.e., the cardinality of V is n. We say that v_i and v_j are neighbors if $(v_i, v_j) \in E$. Typically, many graph mining and learning algorithms can be expressed by computations on vertices of a graph, and the edges of such a graph represents dependency between the data samples stored on the vertices [10, 5].

The goal of GP is to *balance* the parallel work on each machine and *minimize* the communication cost among different machines. In the following, we will briefly introduce two important GP strategies: edge-cut and vertex-cut. Furthermore, we will also discuss the SGP problem based on vertex-cut.

[1]The S in *S-PowerGraph* comes from the term *streaming*.

2.1 GP by Edge-Cut

The traditional GP problem is to balance the number of vertices in each partition and minimize the edges across different partitions, which is known as edge-cut based GP [7, 14, 12, 13, 15]. More specifically, the problem is defined as follows:

$$\min_{A} \ \left|\{e \mid e = (v_i, v_j) \in E, v_i \in V_x, v_j \in V_y, x \neq y\}\right|$$
$$s.t. \ \frac{\max_i |V_i|}{\frac{1}{p}\sum_{i=1}^{p} |V_i|} \leq 1 + \xi,$$

where $|\cdot|$ denotes the cardinality of a set, ξ is a constant with $0 \leq \xi \leq 1$, c is the number of partitions, $A = \{V_1, V_2, \ldots, V_p\}$ is a partition of V, i.e., $V_1 \cup \cdots \cup V_p = V$ and $V_x \cap V_y = \emptyset$ for $x \neq y$.

The objective function in the above formulation tries to minimize the number of edges across different partitions, and the constraint makes the partitions as balanced as possible. The demonstration of edge-cut is shown in Figure 1 (a). Representative methods include the offline method Metis [7] and the streaming methods LDG [14] and FENNEL [15].

2.2 GP by Vertex-Cut

The edge-cut methods are efficient when partitioning sparse graphs, but they are not suitable when adapted for scale-free graphs. Typically, natural graphs are not only large but also skewed. The degree of the graph usually follows the power-law degree distribution:

$$P(d) \propto d^{-\alpha},$$

where $P(d)$ is the probability that a vertex has degree d, α is a positive constant. For example, $\alpha \approx 2.2$ in the twitter-2010 social network [8]. This skewed degree distribution might result in an imbalanced partitioning and challenge the traditional edge-cut methods.

PowerGraph [5] defined a balanced p-way vertex-cut based graph partitioning model. Let each edge $e \in E$ be assigned to a machine $M(e) \in \{1, \ldots, p\}$. Then each vertex spans a set of different partitions $A(v) \subseteq \{1, \ldots, p\}$. Thus, $|A(v)|$ is the number of replications of v among different machines. The balanced vertex-cut based partitioning model is defined as follows:

$$\min_{A} \frac{1}{|V|} \sum_{v \in V} |A(v)|$$
$$s.t. \ \max_{m} |\{e \in E \mid M(e) = m\}| < \lambda \frac{|E|}{p},$$

where $\lambda \geq 1$ is an imbalance factor.

In vertex-cut, the number of edges on each machine is used to estimate the computation cost of that machine, and the number of the replications of the vertices is used to estimate the communication cost. The objective is still to balance the workload among machines and minimize the communication cost.

2.3 SGP by Vertex-Cut

The problem of SGP by vertex-cut tries to optimize the following conditional expectation:

$$\underset{k}{\operatorname{argmin}}\ \mathbb{E}\Big[\sum_{v\in V}|A(v)|\Big|M_i, M(e_{i+1})=k\Big], \qquad (1)$$
$$s.t.\ \max_m|\{e\in E\mid M(e)=m\}|<\lambda\frac{|E|}{p},$$

where $\mathbb{E}$ denotes the expectation operation, k is the index number of a partition, M_i is the assignment of the previous i edges, e_{i+1} is the current edge in the stream.

3. METHODOLOGY

We consider a streaming model in which edges arrive in a stream and each edge appears only once in the stream. The vertices also arrive in a stream order just like what a web crawler does. Moreover, each time a vertex arrives together with some of its edges.

For SGP, the order that the vertex arrives is very important because it can influence the performance of the algorithms significantly [14, 15]. In our streaming setting, the vertices arrive in a particular order with all of their out-edges arriving in a random order. As stated in existing works [14, 15], there are three typical orders in which the vertices arrive:

- **Random (Rnd)**: The vertices simply arrive in a random order or in an order given by a random permutation of the vertices.
- **Breadth-first search (BFS)**: A starting vertex is selected from the graph and new vertices arrive if they are found by a breadth-first search.
- **Depth-first search (DFS)**: This order is just the same as the BFS ordering except that depth-first search will be used.

3.1 Model

We observe that the replication factor, which is defined as the total number of replications divided by the total number of vertices, will be smaller if we cut more vertices with high degree. That is because when the degrees of vertices are under the power-law distribution, there exist many vertices with small degree and a relatively small amount of vertices are dense. So if choices are available, we always choose to cut a vertex with relatively higher degree. Then we can simply use this strategy to improve the naive random (hashed) vertex-cut.

Relative to the model in (1), we add a constraint to guarantee if either vertex of the edge can be cut, the one with higher degree will be chosen. In particular, we defined the following model for streaming partitioning:

$$\underset{k}{\operatorname{argmin}}\ \mathbb{E}\Big[\sum_{v\in V}|A(v)|\Big|M_i, M(e_{i+1})=k, k\in S\Big]. \qquad (2)$$

Here S is a set of partitions which the vertex with higher degree of e_{i+1} belongs to. If neither of the two vertices of e_{i+1} belongs to any partition, we let S simply contain all the k partitions. Thus, we define S as:

$$S=\begin{cases}\{1,...,p\} & \text{if } A(u)=0 \text{ or } A(v)=0\\ A(\hat{w}) & \text{otherwise,}\end{cases}$$

where

$$\hat{w}=\underset{w}{\operatorname{argmin}}\,D(w)\ \text{ s.t. }\ w\in\{u,v\}\ \text{ for }\ (u,v)=e_{i+1}.$$

Here $D(u)$ is the degree of vertex u. In this paper, we simply consider the graph undirected. For particular applications, $D(u)$ can be the degree for undirected graph or directed graph, and it can further be either in-degree or out-degree of u if the graph is directed.

Additionally, since we do the partitioning job in a real online condition, the degree (in-degree and out-degree) of each vertex is not known in advance. The degree distribution is unveiled gradually when the edges of the graph arrive in stream. So we need to find a way to estimate the degree distribution of the whole graph.

We use model in (2) and some kind of estimation of degree distribution to build the algorithm *Random-Degree*, *Degree* and *DegreeIO* in the next section.

3.2 Algorithms

In this paper, we examine multiple partitioning algorithms. We define each of them as follows.

Let $P(k)$ denote the current set of edges belong to the kth partition, and let $[p]=\{1,\ldots,p\}$. We define

$$maxedges=\max_{k\in[p]}\{|P(k)|\}\ \text{ and }\ minedges=\min_{k\in[p]}\{|P(k)|\}.$$

We refer to a certain partition by its index idx. Then the idxth partition is denoted by P_{idx}. For each algorithm, and index idx will be decided by different heuristic and each edge will be assigned to the idxth partition. For the greedy algorithms, when assigning an edge, each algorithm will first compute a *Score* for each partition. It is the partition with the highest *Score* that the edge will finally be assigned to.

The first three algorithms are adapted from some related works.

Random The random method simply assigns each edge via a random hash function as implemented in PowerGraph. So it is the most naive algorithm. The edge e is assigned to P_{idx} where idx is decided by:

$$idx=hash(e)$$

$hash(e)$ here is a randomized hash function.

Grid GraphBuilder [6] provides a Grid-based Constrained Random Vertex-cut algorithm, which can also be adapted for streaming graph partitioning. In this Grid partitioning, a vertex is mapped into a certain shard which contains a constrained set of partitions. Then, the assignment will be chosen from the constrained set randomly. The constraint on the choice of partitions produces less replications than the naive randomized partitioning.
The edge e is assigned to P_{idx} where idx is decided by:

$$idx=Grid_hash(e)$$

The hash function $Grid_hash(e)$ here is the same one implemented by GraphBuilder.

Balance (i.e., PowerGraph) - An enhanced version of the greedy algorithm proposed in PowerGraph [5]. However some particular orders may result in relatively imbalance partitioning for the algorithm introduced by PowerGraph. So in practice the proposed greedy algorithm cannot be used as a baseline. We instead need an algorithm to enhances the constraint of imbalance which turns out to be the algorithm *Balance(PowerGraph)*. A new constraint is added to avoid imbalance. This modified version plays a role as baseline which demonstrates the performance of the PowerGraph partitioning.

$$\begin{aligned} Score(k) &= \mathbf{1}\{k \in A(u)\} + \mathbf{1}\{k \in A(v)\} + balance(k), \\ balance(k) &= \frac{maxedges - |P(k)|}{maxedges - minedges + 1}, \\ (u,v) &= e, \\ k &\in [p]. \end{aligned}$$

where $\mathbf{1}\{\#\} = 1$ if $\#$ is true and $\mathbf{1}\{\#\} = 0$ otherwise.

If $\frac{\max_i |p_i|}{avg|p_i|} \geq 1.1$, then the edge e is assigned to P_{idx} where idx is decided by:

$$idx = \underset{k\in[p]}{\operatorname{argmax}}\{balance(k)\}$$

Otherwise, the edge e is assigned to P_{idx} where idx is decided by:

$$idx = \underset{k\in[p]}{\operatorname{argmax}}\{Score(k)\}$$

The following algorithms are proposed by us which are based on the model in (2).

Random-Degree The edge e is assigned to P_{idx} where idx is decided by:

$$idx = hash(v_L)$$

where

$$v_L = \underset{w}{\operatorname{argmin}}\, D_{in}(w) \;\text{ s.t. }\; w \in \{u, v\}$$

Degree - Assign e to a partition of maximal score, breaking ties randomly. The score of the kth partition is defined as follows:

$$\begin{aligned} Score_{in}(k) &= \mathbf{1}\{k \in A(u)\} + \mathbf{1}\{k \in A(v)\} \\ &\quad + \mathbf{1}\{k \in A(u), D_{in}(u) \leq D_{in}(v)\} \\ &\quad + \mathbf{1}\{k \in A(v), D_{in}(v) \leq D_{in}(u)\} \\ &\quad + balance(k), \\ balance(k) &= \frac{maxedges - |P(k)|}{maxedges - minedges + 1}, \\ (u,v) &= e, \\ k &\in [p]. \end{aligned}$$

If $\frac{\max_i |p_i|}{avg|p_i|} \geq 1.1$, then the edge e is assigned to P_{idx} where idx is decided by:

$$idx = \underset{k\in[p]}{\operatorname{argmax}}\{balance(k)\}$$

Otherwise, the edge e is assigned to P_{idx} where idx is decided by:

$$idx = \underset{k\in[p]}{\operatorname{argmax}}\{Score_{in}\}$$

DegreeIO - The score of the kth partition is defined as follows:

$$\begin{aligned} Score_{in}(k) &= \mathbf{1}\{k \in A(u)\} + \mathbf{1}\{k \in A(v)\} \\ &\quad + \mathbf{1}\{k \in A(u), D_{in}(u) \leq D_{in}(v)\} \\ &\quad + \mathbf{1}\{k \in A(v), D_{in}(v) \leq D_{in}(u)\} \\ &\quad + balance(k) \\ Score(k) &= \mathbf{1}\{k \in A(u)\} + \mathbf{1}\{k \in A(v)\} \\ &\quad + \mathbf{1}\{k \in A(u), D(u) \leq D(v)\} \\ &\quad + \mathbf{1}\{k \in A(v), D(v) \leq D(u)\} \\ &\quad + balance(k), \\ balance(k) &= \frac{maxedges - |P(k)|}{maxedges - minedges + 1}, \\ (u,v) &= e, \\ k &\in [p]. \end{aligned}$$

For the current edge e, we let u be the source vertex and v be the target vertex. If $D_{out}(u) \geq p$ and $D_{out}(v)$ is unknown which means vertex v has not arrived in the stream, then we do not assign the edge immediately and store it in a buffer until vertex v arrives. If $D_{out}(v)$ is known, the edge e is assigned to P_{idx} where idx is decided by:

$$idx = \underset{k\in[p]}{\operatorname{argmax}}\{Score(k)\}$$

Otherwise, the edge e is assigned to P_{idx} where idx is decided by:

$$idx = \underset{k\in[p]}{\operatorname{argmax}}\{Score_{in}(k)\}$$

3.3 Theoretical Analysis

Different algorithms have different motivations.

The first three algorithms are from related works and relatively need less computation and memory.

For all the greedy algorithms *Balance(Powergraph)*, *Degree* and *DegreeIO*, when the graph is not very skewed, *Balance* is a good choice which requires less computation and memory.

Random-Degree demonstrates the main idea of our model: decisions of which partition to be assigned to is always according to the vertex with lower degree. Although this algorithm cannot be used in the stream setting because the degree of each vertex is unknown in a stream, a theoretical analysis is available to show how our idea reduce the replication factor.

THEOREM 1. *A random-degree vertex-cut on p machines has the expected replication factor:*

$$\begin{aligned} &\mathbb{E}\left[\frac{1}{|V|}\sum_{v\in V}|A(v)|\right] \\ &= 1 + \frac{(p-1)}{|V|}\sum_{v\in V}\left(1 - (1-\frac{1}{p})^{(1-Ratio(v))D[v]}\right) \end{aligned}$$

Here $D(v)$ is the degree of vertex v. And all edges are treated as undirected. $Ratio(v)$ is a ratio of degree of each vertex.

PROOF. Our degree-random vertex-cut algorithm: for edge $e(u,v)$, $v_L = arg\min_w\{D(w)|w \in \{u,v\}\}$ which is the vertex with lower degree. Then we assign the edge by some kind of hash function. So we have the assignment of e:

$$A(e) = hash(v_L),$$

Denote the indicator P_i as the event that vertex v has at least one of $S(v)$ in the ith partition where $S(v)$ is a subset of the adjacent edges of vertex v. Then follow the above procedure, the expectation $\mathbb{E}[P_i]$ is:

$$\mathbb{E}[P_i] = 1 - (1 - \frac{1}{p})^{D_s[v]}$$

where $D_s(v)$ is the number of edges in $S(v)$.

We denote the fraction of the adjacent edges of v that satisfy $D(v) \leq D(u)$ as $Ratio(v)$ where u is the another vertex of each adjacent edge $e(u,v)$.

Intuitively, if we assume $P(d) = Cd^{-\alpha}$ I have:

$$1 - Ratio(v) = P\left(d \leq D(v)\right) = \sum_{d=1}^{D(v)} P(d) = C\sum_{d=1}^{D(v)} d^{-\alpha}$$

For some vertex v, $Ratio(v)$ of its adjacent edges will be hashed to one certain partition, and the residual will be randomly hashed to all the p partitions. So for the $Ratio(v)$ of the adjacent edges, the number of replication of v is simply 1.

And for the residual $1 - Ratio(v)$ of the adjacent edges, we have $\mathbb{E}[P_i] = 1-(1-\frac{1}{p})^{(1-Ratio(v))D[v]}$. Here P_i will the other $p-1$ partitions than the one that already have a replication.

Put the two parts together, then we can have:

$$\begin{aligned}
&\mathbb{E}\left[|A(v)|\right] \\
&= 1 + \sum_{i=1}^{p-1} \mathbb{E}[P_i] \\
&= 1 + (p-1)\left(1 - (1 - \frac{1}{p})^{(1-Ratio(v))D[v]}\right)
\end{aligned}$$

Thus the expected replication factor is:

$$\begin{aligned}
&\mathbb{E}\left[\frac{1}{|V|}\sum_{v\in V}|A(v)|\right] \\
&= \frac{1}{|V|}\sum_{v\in V}\left(1 + (p-1)\left(1 - (1 - \frac{1}{p})^{(1-Ratio(v))D[v]}\right)\right) \\
&= 1 + \frac{(p-1)}{|V|}\sum_{v\in V}\left(1 - (1 - \frac{1}{p})^{(1-Ratio(v))D[v]}\right)
\end{aligned}$$

□

The two novel greedy algorithms *Degree* and *DegreeIO* can be viewed as a greedy version of *Random-Degree* in the stream setting.

Degree is a greedy procedure which makes use of the in-degree distribution. The idea seems similar to that proposed in PowerLyra [3] which acts as a component of the distributed framework. However, the motivations of both are different. Unlike PowerLyra, we do not assume that the application adapted for the graph is known; i.e., we do not know which kind of edges (in-edges or out-edges) the computation involves in advance. Thus, we treat each edge as undirected. But the current problem is that the out-degree of a vertex is unknown until it arrives with all its out-edges. Fortunately, the degree distribution of the in-edges or the relationship of the degrees between vertices can be easily estimated by the vertices arrived according to our streaming model. This is why we only consider the in-degree in this step.

DegreeIO aims to meet the challenge that both in-degree and out-degree of natural graphs are skewed and to further improve the vertex-cut. To solve the problem that the out-degree of some vertex is unknown temporarily, the edge is buffered until the out-degree is known. For the vertex whose out-degree is small, we simply reuse the *Degree* step to reduce the memory use.

So if we are sure in advance that the in-degree of the graph is more skewed than the out-degree, *Degree* is a better choice with less computation and memory use. Otherwise *DegreeIO* is always a better choice with slightly more computation cost. The *Degree* and *DegreeIO* are two variants of our S-PowerGraph method.

4. EXPERIMENT

In this section we demonstrate the setup and results of our experiments as well as the discussion of the results.

We use a server with Intel Xeon cores, 64GB of RAM.

4.1 Datasets

The collection of large graph datasets we use in the experiments consists of real-world graphs shown in Table 1(a) and synthetic power-law graphs shown in Table 2. Some of the real-world graphs were selected to be the same as the experiment of PowerLyra[3]. And some additional real-world graphs were from the UF Sparse Matrices Collection[4]. The synthetic power-law graphs were generated by a combination of two synthetic graphs. For one of them, the out-degree of each vertex is sampled randomly from a Zipf distribution[1] with parameter α and the in-degree of each vertex is under a uniform distribution. For the other one, the in-degree of each vertex is sampled randomly from a Zipf distribution with parameter β and the out-degree of each vertex is under a uniform distribution. The power-law constant (α and β) of the synthetic graphs ranges from 1.9 to 2.2. The smaller the α is, the skewer the graph will be. However the different power-law constant result in different partitioning performance. And the relationship between the distribution of in-degree and out-degree also acts as an important role. To demonstrate such fact, we separate the synthetic graph collection into two parts: the one with parameter α no smaller

than β shown in Table 2 (a) and the one with parameter β greater than α shown in Table 2 (b).

Table 1: Real-world graphs

Alias	Graph	$\|V\|$	$\|E\|$
Tw	Twitter[8]	42M	1.47B
It	It-2004[4]	41M	1.15B
UK	UK-2005[2]	40M	936M
Arab	Arabic-2005[4]	22M	0.6B
Wiki	Wiki[2]	5.7M	130M
LJ	LiveJournal[11]	5.4M	79M
HW	Hollywood-2009[4]	1.1M	57M
In	In-2004[4]	1.4M	17M
As	As-skitter[4]	1.7M	11M
CP	CoPapersCiteseer[4]	0.4M	16M
WG	WebGoogle[9]	0.9M	5.1M

Table 2: Synthetic graphs with different power-law constants

(a) Synthetic graphs with $\alpha \geq \beta$

Alias	α	β	#Edges
S1	2.2	2.2	71,334,974
S2	2.2	2.1	88,305,754
S3	2.2	2.0	134,881,233
S4	2.2	1.9	273,569,812
S5	2.1	2.1	103,838,645
S6	2.1	2.0	164,602,848
S7	2.1	1.9	280,516,909
S8	2.0	2.0	208,555,632
S9	2.0	1.9	310,763,862

(b) Synthetic graphs with $\alpha < \beta$

Alias	α	β	#Edges
S10	2.1	2.2	88,617,300
S11	2.0	2.2	135,998,503
S12	2.0	2.1	145,307,486
S13	1.9	2.2	280,090,594
S14	1.9	2.1	289,002,621
S15	1.9	2.0	327,718,498

4.2 Baseline and Evaluation Metric

4.2.1 Baseline

Table 3 lists several representative graph partitioning methods and our S-PowerGraph method. Note that the *Degree* and *DegreeIO* are two variants of our S-PowerGraph. In our experiment, we take algorithm *Random*, *Grid*, and *Balance(PowerGraph)* as the baseline. Note that these algorithms comes from PowerGraph and GraphBuilder listed in Table 3. The reason why the rest of them are not considered is that the first two are of edge-cut which cannot be compared to the results of vertex-cut and the algorithm of PowerLyra cannot be used in the stream setting.

4.2.2 Evaluation Metric

The result of the experiments can be described by three factors: communication cost, imbalance and sensibility to number of partitions. We used the replication factor to define the communication cost between different partitions which was introduced by PowerGraph[5]:

$$\lambda = \frac{1}{|V|} \sum_{v \in V} |A(v)|$$

, where V is the set of all the vertices, $|V|$ denotes the number of vertices and $|A(v)|$ denotes the number of mirrors (replications) of vertex v. We experiment the replication factor on all the real-world graphs and synthetic graphs.

To make the result more clear, wee also illustrate the improvement of our algorithms over the baseline. In all the three algorithms which act as the baseline, the algorithm *Balance(PowerGraph)* enjoys the lowest replication factor defined above. So we compare the replication factor of our algorithms to that of *Balance(PowerGraph)* and denote the improvement as:

$$improvement = \frac{\lambda_{PowerGraph} - \lambda_{Algorithm}}{\lambda_{PowerGraph}} \times 100\%$$

where $\lambda_{Algorithm}$ is denoted by the replication factor produced by a certain algorithm.

The second factor is the imbalance factor defined as:

$$\rho = \frac{\max_m |\{e \in E \mid A(e) = m\}|}{|E|/p}$$

, where $|A(e)|$ denotes the partition that edge e is placed in , $|E|$ is the number of edges of the graph and p is the number of partitions. However, by using the vertex-cut strategy, the imbalance is no longer a big deal. So we just show the results of the real-world graphs.

Finally we tested the sensibility to the number of partitions of our algorithm. We used the twitter-2010 social network[8] and compared the replication factor by the increasing number of partitions. And we also show that the algorithm is scalable with the power-law constant by taking experiments on synthetic graphs with different power-law constant.

4.3 Experimental Results

In Figure 2 (a),(b) and (c), we compared the replication factor on different real-world graphs in different orders by different algorithms. In all the experiments the number of partitions is 48. Our algorithm *DegreeIO* produced lower replication factor in each experiment. Figure 2 (d),(e) and (f) shows the improvement of the algorithms *Degree* and *DegreeIO* compared to *Balance(PowerGraph)*.

In Figure 3 (a), (b), (c), (d), (e) and (f), we compared the replication factor on synthetic graphs with different power-law constants in different orders by different algorithms. In all the experiments the number of partitions is 48. Figure 4 (a), (b), (c), (d), (e) and (f) shows the improvement of the algorithms *Degree* and *DegreeIO* compared to *Balance(PowerGraph)*.

In Table 4, Table 5 and Table 6, we compared the imbalance factor on different real-world graphs in different orders by

Table 3: Some representative graph partitioning methods.

Method	LDG	FENNEL	PowerGraph	PowerLyra	GraphBuilder	S-PowerGraph
Partition Strategy	edge-cut	edge-cut	vertex-cut	hybrid edge-cut /vertex-cut	vertex-cut	vertex-cut
Balance Type	vertex-balance	vertex-balance	edge-balance	edge-balance	edge-balance	edge-balance
Streaming/Online	yes	yes	no	no	no	yes

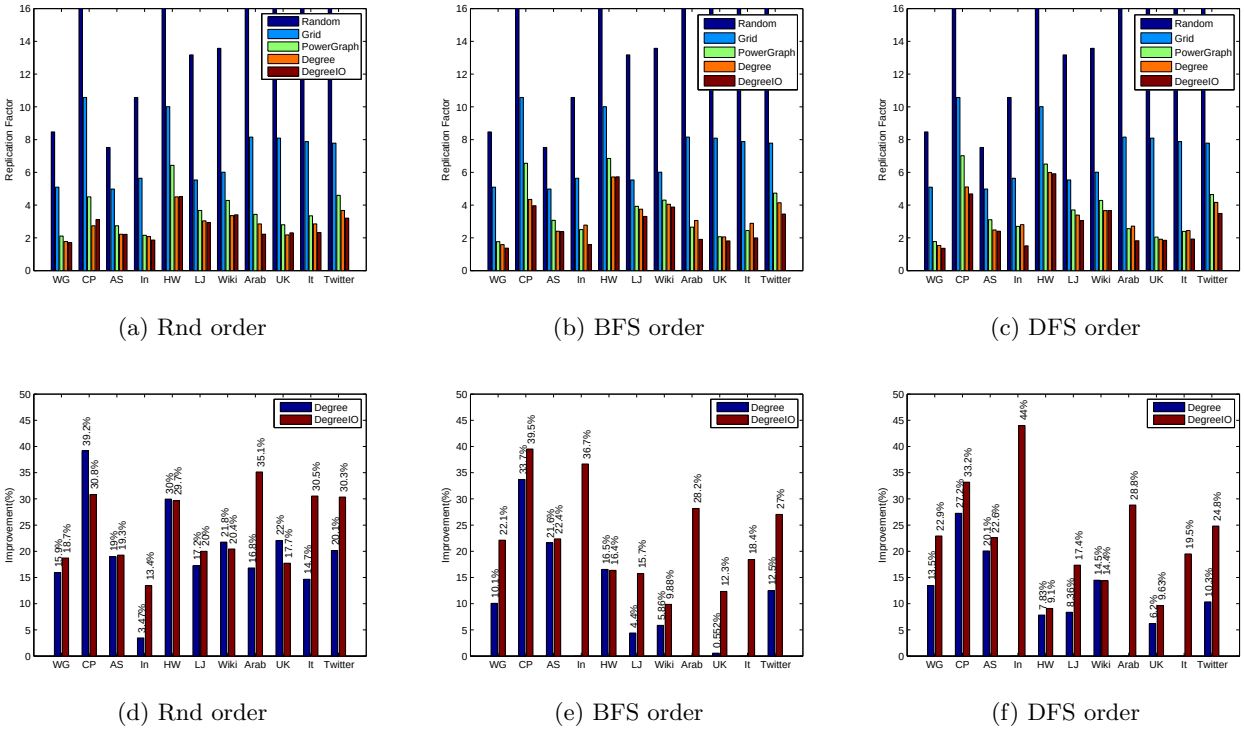

(a) Rnd order (b) BFS order (c) DFS order

(d) Rnd order (e) BFS order (f) DFS order

Figure 2: Replication factor (first row) and improvement (second row) on real-world graphs in different stream orders.

different algorithms. In all the experiments the number of partitions is 48. The imbalance factors of our algorithms are a little bit higher but still tiny. Most of them are smaller than 1.01.

Table 4: Imbalance factor on real-world graphs in random order (Here PowerGraph is abbreviated to PowerG. for a better layout.)

	Method				
Graph	Random	Grid	PowerG.	Degree	DegreeIO
WG	1.00725	1.04038	1.00024	1.00044	1.00041
CP	1.00341	1.01943	1.00015	1.0005	1.00122
As	1.00408	1.06652	1.00005	1.00024	1.00018
In	1.00394	1.02741	1.00075	1.01349	1.00298
HW	1.00227	1.02349	1.00002	1.00008	1.00008
LJ	1.00173	1.00703	1.00005	1.00006	1.00005
Wiki	1.00145	1.03101	1.00002	1.0001	1.00003
Arab	1.00057	1.01618	1.00007	1.00055	1.00006
UK	1.0005	1.02545	1.00002	1.00144	1.00003
It	1.00043	1.01874	1.00003	1.00031	1.00005
Tw	1.00032	1.04214	1	1	1

In Figure 5 (a), (b) and (c), we evaluate our algorithm on

Table 5: Imbalance factor on real-world graphs in BFS order (Here PowerGraph is abbreviated to PowerG. for a better layout.)

	Method				
Graph	Random	Grid	PowerG.	Degree	DegreeIO
WG	1.00725	1.04038	1.00015	1.00036	1.10936
CP	1.00341	1.01943	1.00007	1.00032	1.01744
As	1.00408	1.06652	1.00003	1.00021	1.00022
In	1.00394	1.02741	1.01396	1.0163	1.04447
HW	1.00227	1.02349	1.0245	1.0399	1.00011
LJ	1.00173	1.00703	1.00001	1.00721	1.00006
Wiki	1.00145	1.03101	1	1.00134	1.00001
Arab	1.00057	1.01618	1.0001	1.00462	1.00848
UK	1.0005	1.02545	1.00046	1.00069	1.00002
It	1.00043	1.01874	1.00024	1.01175	1.06839
Tw	1.00032	1.04214	1	1	1

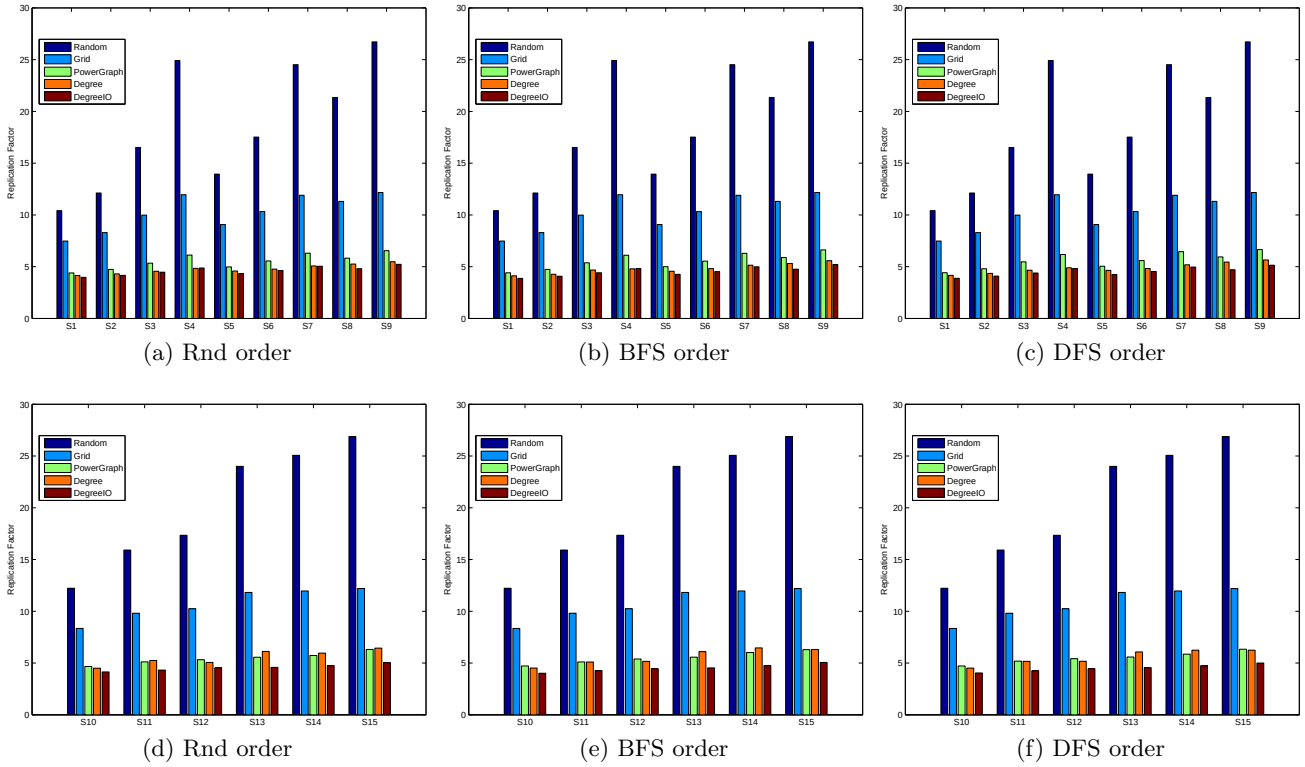

(a) Rnd order (b) BFS order (c) DFS order

(d) Rnd order (e) BFS order (f) DFS order

Figure 3: Replication factor on synthetic graphs in different stream orders.

Table 6: Imbalance factor on real-world graphs in DFS order (Here PowerGraph is abbreviated to PowerG. for a better layout.)

	Method				
Graph	Random	Grid	PowerG.	Degree	DegreeIO
WG	1.00725	1.04038	1.00014	1.00043	1.00024
CP	1.00341	1.01943	1.00009	1.00027	1.00035
As	1.00408	1.06652	1.00002	1.00015	1.00019
In	1.00394	1.02741	1.0161	1.01655	1.04168
HW	1.00227	1.02349	1.03225	1.04545	1.00177
LJ	1.00173	1.00703	1.00003	1.00534	1.00007
Wiki	1.00145	1.03101	1	1.00002	1.00001
Arab	1.00057	1.01618	1.0003	1.00787	1.02141
UK	1.0005	1.02545	1.00022	1.00568	1.00002
It	1.00043	1.01874	1.00018	1.047	1.00005
Tw	1.00032	1.04214	1	1	1

Twitter-2010 in different order to compare the sensibility to the number of partitions. Our algorithms is always better and the improvement gets more obvious when the number of partitions increases.

4.4 Discussion

Taking into account all the experiments, our algorithm *DegreeIO* is better than *Balance(PowerGraph)* in all datasets. And *Degree* improves the performance of *Balance(PowerGraph)* in most cases. For the real-world graphs, we can see that the improvement on WebGoogle is not very obvious. The reason may be that the WebGoogle graph is not so skewed and the power-law constant may be high. For the largest dataset Twitter-2010, *Degree* gets an improvement for about 15% to 20% while *DegreeIO* gets nearly 25% to 30%. Our algorithms produce more improvement than the algorithm introduced by PowerGraph[5] when the graphs are more skewed in the degree distribution which can be shown in the experiment on synthetic graphs. This kind of phenomenon is a consequence of the motivation of our algorithm. Recall that our algorithm makes use of the power-law degree distribution to further reduce the vertex-cut. So our algorithm will not produce an obvious improvement if the distribution of degrees are nearly uniform, and it works better when the graph is more skewed. Finally the experiments on Twitter-2010 shows good scalability for our algorithms.

Note that the algorithm *Degree* does not always produce a better replication factor than *Balance(PowerGraph)*. The reason is that *Degree* estimates the degree distribution only by the in-degree. Thus if the out-degree is more skewed than the in-degree, for example In-2004, Arabic-2005, It-2004 and the second part of the synthetic graph collection, the estimation will be wrong which will result in bad partitioning. And *DegreeIO* estimate the degree distribution by both in-degree and out-degree, so the performance is better in such datasets. For example, the improvement of Arabic-2005 is nearly 35%.

5. CONCLUSION

In this paper we have studied streaming graph partitioning by vertex-cut. In particular, we have proposed a novel streaming graph partitioning method for natural graphs by adopting vertex-cut strategy and called it *S-PowerGraph*. To the best of our knowledge, this is the first work to sys-

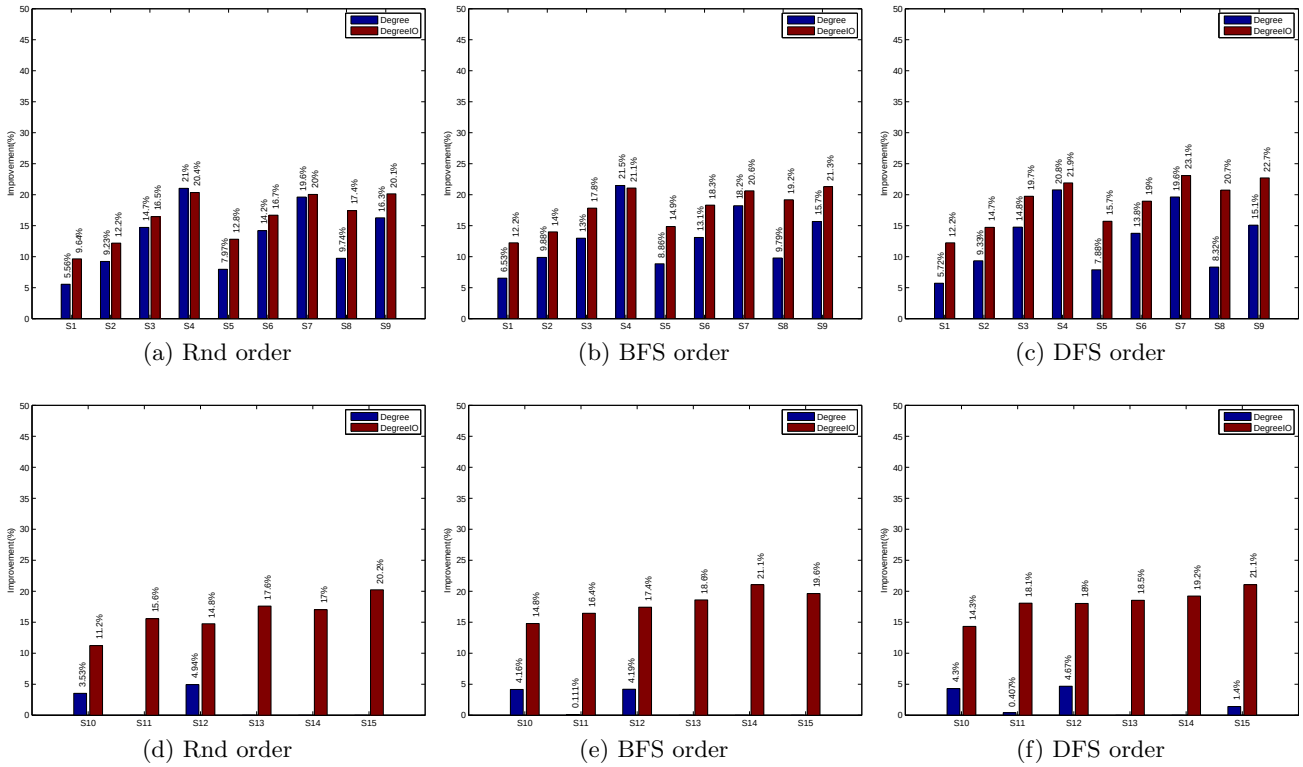

(a) Rnd order (b) BFS order (c) DFS order

(d) Rnd order (e) BFS order (f) DFS order

Figure 4: Improvement on synthetic graphs in different stream orders (compared with PowerGraph).

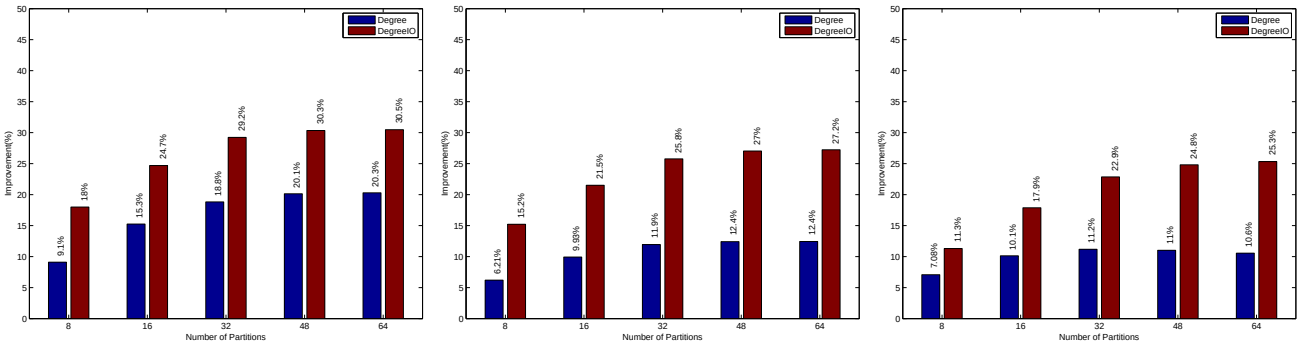

(a) Rnd order (b) BFS order (c) DFS order

Figure 5: Improvement on Twitter-2010 in different stream orders (compared with PowerGraph).

tematically study the effectiveness of vertex-cut in streaming graph partitioning. Experiments on large graphs show that our method is more suitable for partitioning skewed natural graphs than the previous approaches with an acceptable imbalance factor. Our method can be used to partition the raw graph data from a web crawler or the graph data loaded from a single storage device. In our future work, we would like to pursued more variants of the proposed strategy theoretically and empirically.

6. REFERENCES

[1] L. A. Adamic and B. A. Huberman. Zipf's law and the internet. *Glottometrics*, 3(1):143–150, 2002.

[2] P. Boldi and S. Vigna. The webgraph framework I: compression techniques. In *Proceedings of the 13th international conference on World Wide Web*, pages 595–602, 2004.

[3] R. Chen, J. Shi, Y. Chen, H. Guan, and H. Chen. Powerlyra: Differentiated graph computation and partitioning on skewed graphs. Technical Report IPADSTR-2013-001, Institute of Parallel and Distributed Systems, Shanghai Jiao Tong University, November 2013.

[4] T. A. Davis and Y. Hu. The University of Florida sparse matrix collection. *ACM Transactions on Mathematical Software*, 38(1):1, 2011.

[5] J. E. Gonzalez, Y. Low, H. Gu, D. Bickson, and C. Guestrin. Powergraph: Distributed graph-parallel

computation on natural graphs. In *Proceedings of the 10th USENIX Symposium on Operating Systems Design and Implementation (OSDI)*, pages 17–30, 2012.

[6] N. Jain, G. Liao, and T. L. Willke. Graphbuilder: scalable graph etl framework. In *First International Workshop on Graph Data Management Experiences and Systems*, page 4, 2013.

[7] G. Karypis and V. Kumar. Multilevel graph partitioning schemes. In *Proceedings of the International Conference on Parallel Processing*, pages 113–122, 1995.

[8] H. Kwak, C. Lee, H. Park, and S. Moon. What is twitter, a social network or a news media. In *Proceedings of the 19th international conference on World Wide Web*, pages 591–600, 2010.

[9] J. Leskovec. Stanford large network dataset collection. *URL http://snap. stanford. edu/data/index. html*, 2011.

[10] Y. Low, J. Gonzalez, A. Kyrola, D. Bickson, C. Guestrin, and J. M. Hellerstein. GraphLab: A new framework for parallel machine learning. In *Proceedings of the Conference on Uncertainty in Artificial Intelligence (UAI)*, pages 340–349, 2010.

[11] A. Mislove, M. Marcon, K. P. Gummadi, P. Druschel, and B. Bhattacharjee. Measurement and analysis of online social networks. In *Proceedings of the 7th ACM SIGCOMM conference on Internet measurement*, pages 29–42, 2007.

[12] J. Nishimura and J. Ugander. Restreaming graph partitioning: simple versatile algorithms for advanced balancing. In *Proceedings of the ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 1106–1114, 2013.

[13] I. Stanton. Streaming balanced graph partitioning algorithms for random graphs. In *Proceedings of the ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1287–1301, 2014.

[14] I. Stanton and G. Kliot. Streaming graph partitioning for large distributed graphs. In *Proceedings of the ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 1222–1230, 2012.

[15] C. E. Tsourakakis, C. Gkantsidis, B. Radunovic, and M. Vojnovic. FENNEL: Streaming graph partitioning for massive scale graphs. In *Proceedings of the ACM International Conference on Web Search and Data Mining*, 2014.